SearchSMBStorage.com » FAQ

Using clustered network-attached storage (NAS) to manage unstructured data

The growth of structured and unstructured data continues to grow at an explosive rate in most storage environments, resulting in a constantly expanding data footprint. Fortunately, clustered storage is becoming a common way to combat this problem. Take this opportunity to download this brief expert guide to learn all about clustered network-attached storage (NAS) as a solution, and the vendors in the NAS market.

TechTarget
Where Serious Technology Buyers Decide

Using clustered network-attached storage (NAS) to manage unstructured data

By: Marc Staimer

Marc Staimer, president of Dragon Slayer Consulting, discusses the pros and cons of clustered network-attached storage (NAS) and the vendors in the NAS market in this FAQ.

What is clustered NAS and how does it differ from traditional NAS?

Clustered NAS, or network-attached storage, is different from typical network-attached storage in that it uses a distributed file system that runs simultaneously on multiple nodes or servers. The key differentiating factor between clustered NAS and traditional NAS is its ability to stripe the data and metadata across the storage nodes and subsystems. This provides access to the network file system from any of the clustered nodes unrelated to the actual location of the data.

What are the benefits of clustering NAS?

It's a scaling situation. There's an old saying with network-attached storage that goes something like this: "I loved my first NAS. I really liked my second. By my tenth I was pulling my hair out."

As network-attached storage scales you get to a point where you can't put anything more in a filer, so you have to get another filer. Then you have to manage both filers -- so you think "OK, so I have two times the workload." Well, not exactly. You are going to have more than two times the workload because you have to start moving data between them and make sure you are getting the optimum use of both. That being said, two still isn't that bad. Three? With three, the balancing gets a lot tougher. Once you get up to 10, it's a nightmare.

So, a clustered NAS solution is designed to provide a single image, a single mount point, and it does all of that load balancing among the nodes without any human intervention. So, you have, in effect, one very large bucket, and you can add more nodes as you go.

What are the drawbacks of clustered network-attached storage?

The biggest drawback of clustered NAS is that most clustered NAS products are designed for a lot of data streams and a lot of users. A single data stream for performance isn't going to perform better, and in some cases, it may perform worse than a single data stream in traditional NAS.

Who needs clustered network-attached storage? Is there a typical use case?

There are a lot of use cases. Unstructured data today is growing much faster than structured data. PowerPoint presentations, MP3s, Microsoft Word documents, the things that people use on their laptops and desktops -- all of this stuff is growing much faster than structured information such as databases, enterprise resource planning (ERP) and mail servers. In fact, this year, unstructured data will surpass structured data in most data centers.

So, just about everyone needs filer or NAS-type systems. The clustered NAS products have a nice leg up in an environment where you have lots of users accessing the same files. So, for example, in the verticals, they have been doing this for a while. The entertainment and music industry uses clustered NAS quite a bit, because it allows you to share workflows.

Think of that from a video or film production point of view, and that's very nice. And, it scales. Easier than, let's say, a storage area network (SAN) file system. Life sciences, pharmaceuticals, oil and gas -- these have been the industries using clustered NAS. Now, having said that, it's becoming more mainstream. One area where you are going to see it grow considerably is in the cloud-based data storage offerings.

Who are the major vendors in the clustered network-attached storage market?

There are quite a few. In no particular order:

- ☐ LustreFS
- ☐ Exanet
- ☐ Ibrix
- ☐ Hewlett-Packard (HP) Enhanced File Services (formerly PolyServe)
- ☐ GlusterFS
- ☐ Panasas
- ☐ Bycast
- ☐ NetApp Ontap GX
- ☐ Isilon Systems
- ☐ IBM Scale-out File Services
- ☐ Active Circle
- ☐ Permabit
- ☐ BlueArc